

Differentiable Parsing and Visual Grounding of Natural Language Instructions for Object Placement

Zirui Zhao, Wee Sun Lee, and David Hsu

Abstract—We present a new method, PARsing And visual GrOUnding (PARAGON), for grounding natural language in object placement tasks. Natural language generally describes objects and spatial relations with *compositionality* and *ambiguity*, two major obstacles to effective language grounding. For compositionality, PARAGON parses a language instruction into an object-centric graph representation to ground objects individually. For ambiguity, PARAGON uses a novel particle-based graph neural network to reason about object placements with uncertainty. Essentially, PARAGON integrates a parsing algorithm into a probabilistic, data-driven learning framework. It is fully differentiable and trained end-to-end from data for robustness against complex, ambiguous language input.

I. INTRODUCTION

Robot tasks, such as navigation, manipulation, assembly, . . . , often involve spatial relations among objects. To carry out tasks instructed by humans, robots must understand natural language instructions about objects and their spatial relations. This work focuses specifically on the object placement task with language instructions. Humans provide verbal instructions to robots to pick up an object and place it in a specific location. The robot must generate object placements based on both language description and visual observation. However, the language expressions about spatial relations are generally ambiguous and compositional, two major obstacles to effective language grounding.

Two types of ambiguity, *positional ambiguity* and *referential ambiguity*, are focused upon in this study. *Positional ambiguity* occurs when humans describe directional relations without specifying exact distances (e.g., to the left side of a plate”). Reference objects are often required to describe spatial relations. When placing an object next to a reference object without specifying its distances, it is hard to link the reference expressions to the referred objects to learn visual grounding. *Referential ambiguity* arises when descriptions of objects are ambiguous, leading to a reference expression being grounded to multiple semantically identical objects and resulting in a multimodal distribution of correct placement.

The compositional structure of language-described spatial relations stems from the visual scene’s and natural language’s compositional nature. A complex scene includes multiple basic objects, and to describe the desired state of a complex scene, one can compose many simple sentences for referents

The authors are from National University of Singapore. Emails: {zirui, leews, dyhsu}@comp.nus.edu.sg. This research is supported in part by the National Research Foundation (NRF), Singapore and DSO National Laboratories under the AI Singapore Program (AISG Award No. AISG2-RP-2020-016) and the Agency of Science, Technology and Research, Singapore, under the National Robotics Program (Grant No. 192 25 00054).

The data, code and appendix are available at <https://bit.ly/ParaGonProj>.




Fig. 1. PARAGON takes as input a language instruction and an image of the task environment. It outputs candidate placements for a target object. The presence of multiple semantically identical objects (e.g., “silver mug”) and omitted distance information cause difficulty for placement generation, and the compositional instructions increase the data required for learning.

and their relations to form a complex language sentence (e.g., the instruction in Fig 1). This characteristic increases the data required for learning compositional language instructions.

To address the issues, we propose PARAGON, a PARsing And visual GrOUnding method for language-conditioned object placement. The core idea of PARAGON is to parse human language into object-centric relations for visual grounding and placement generation, and encode those procedures in neural networks for end-to-end training. The parsing module of PARAGON decomposes compositional instructions into object-centric relations, enabling the grounding of objects separately without compositionality issues. The relations can be composed and encoded in graph neural networks (GNN) for placement generation. This GNN uses particle-based message-passing to model the uncertainty caused by ambiguous instructions. All the modules are encoded into neural networks and connected for end-to-end training, avoiding the need for individual training module labeling.

PARAGON essentially integrates parsing-based methods into a probabilistic, data-driven framework, exhibiting robustness from its data-driven property, as well as generalizability and data efficiency due to its parsing-based nature. It also adapts to the uncertainty of ambiguous instructions using particle-based probabilistic techniques. The experiments demonstrate that PARAGON outperforms the state-of-the-art method in language-conditioned object placement tasks in the presence of ambiguity and compositionality.

II. RELATED WORK

Grounding human language in robot instruction-following has been studied in the recent decades [1], [2], [3], [4], [5], [6], [7], [8], [9]. Our research focuses on object placement instructed by human language. In contrast to picking [1], [2] that needs only a discriminative model to ground objects from reference expressions, placing [3], [4], [5], [6], [7] requires a generative model conditioned on the relational constraints of object placement. It requires capturing com-

plex relations between objects in natural language, grounding reference expression of objects, and generating placement that satisfies the relational constraints in the instructions.

Parsing-based methods for robot instruction-following [3], [4], [6], [10], [11], [12], [13], [14] parse natural language into formal representations using hand-crafted rules and grammatical structures. Those hand-crafted rules are generalizable but not robust to noisy language [15]. Among these studies, those focusing on placing [3], [4], [6] lack a decomposition mechanism for compositional instructions and assume perfect object grounding without referential ambiguity. Recently, [8], [9], [7] used sentence embeddings to learn a language-condition policy for robot instruction following, which are not data-efficient and hard to generalize to unseen compositional instructions. We follow [16] to integrate parsing-based modules into a data-driven framework. It is robust, data-efficient, and generalizable for learning compositional instructions. We also use probabilistic techniques to adapt to the uncertainty of ambiguous instructions.

PARAGON has a GNN for relational reasoning and placement generation, which encodes a mean-field inference algorithm similar to [17]. Moreover, our GNN uses particles for message passing to capture complex and multimodal distribution. It approximates a distribution as a set of particles [18], which provides strong expressiveness for complex and multimodal distribution. It was used in robot perception [19], [20], [21], recurrent neural networks [22], and graphical models [23], [24]. Our approach employed this idea in GNN for particle-based message passing.

III. OVERVIEW

We focus on the grounding human language for tabletop object placing tasks. In this task, scenes are composed of a finite set of 3D objects on a 2D tabletop. Humans give natural language instruction $\ell \in \mathcal{L}$ to guide the robot to pick an object and put it at the desired position $\mathbf{x}_{\text{tgt}}^*$. The human instruction is denoted as a sequence $\ell = \{\omega_l\}_{1 \leq l \leq L}$ where ω_l is a word, e.g., instruction in Fig.6 is $\{\omega_1 = \text{put}, \omega_2 = \text{a}, \dots\}$. An instruction should contain a target object expression (e.g., “a plate”) to specify the object to pick and express at least one spatial relation (e.g., “next to a silver mug”) for placement description. The robot needs to find the target object’s placement distribution $p(\mathbf{x}_{\text{tgt}}^* | \ell, \mathbf{o})$ conditioned on the language instruction ℓ and visual observation \mathbf{o} .

The core idea in our proposed solution, PARAGON, is to leverage object-centric relations from linguistic and visual inputs to perform relational reasoning and placement generation, and encode those procedures in neural networks for end-to-end training. The pipeline of PARAGON is in Fig 2. PARAGON first uses the soft parsing module to convert language inputs “softly” into a set of relations, represented as triplets. A grounding module then aligns the mentioned objects in triplets with objects in the visual scenes. The triplets can form a graph by taking the objects as the nodes and relations as the edges. The resulting graph is fed into a GNN for relational reasoning and generating placements. The GNN encodes a mean-field inference algorithm for a

conditional random field depicting spatial relations in triplets. PARAGON is trained end-to-end to achieve the best overall performance for object placing without annotating parsing and object-grounding labels.

IV. SOFT PARSING

The soft parsing module is to extract spatial relations in complex instructions for accurate placement generation. The pipeline is in Fig 4. Dependency trees capture the relations between words in natural language, which implicitly indicate the relations between the semantics those words express [25], [26]. Thus, we use a data-driven approach to explore the underlying semantic relations in the dependency tree for extracting relations represented as relational triplets. It takes linguistic input and outputs relational triplets, where the triplets’ components are represented as embeddings.

A. Preliminaries

Triplets: A triplet consists of two entities and their relation, representing a binary relation. Triplet provides a formal representation of knowledge expressed in natural language, which is widely applied in scene graph parsing [25], relation extraction [27], and knowledge graph [28]. The underlying assumption of representing natural language as triplets is that natural language rarely has higher-order relations, as humans mostly use binary relations in natural language [29]. For spatial relations, two triplets can represent ternary relations (e.g., “between A and B” equals “the right of A and left of B” sometimes). As such, it is sufficient to represent instructions as triplets for common object-placing purposes.

Dependency Tree: A dependency tree (shown in Fig. 3) is a universal structure that examines the relationships between words in a phrase to determine its grammatical structure [30]. It uses part-of-speech tags to mark each word and dependency tags to mark the relations between two words. A part-of-speech tag [31] categorizes words’ correspondence with a particular part of speech, depending on the word’s definition and context, such as in Fig. 3, “cup” is a “Noun”. Dependency tags mark two words relations in grammar, represented as Universal Dependency Relations [32]. For example, in Fig. 3, the Noun “cup” is the “direct object (Dobj)” of the Verb “place”. Those relations are universal. A proper dependency tree relies on grammatically correct instructions, whereas noisy language sentences may result in imperfect dependency trees. Thus, we use the data-driven method to adapt to imperfect dependency trees.

B. Method

To make the parsing differentiable, the soft parsing module “softens” the triplet as the attention to the words of linguistic inputs $\{a_{\gamma,l}\}_{1 \leq l \leq L}, \sum_l a_{\gamma,l} = 1, \gamma \in \{\text{subj}, \text{obj}, \text{rel}\}$. We compute the embeddings of components in triplets by the attention-weighted sum of the word embeddings $\varphi_\gamma = \sum_l a_{\gamma,l} f_{\text{CLIP}}(\omega_l)$. The word embeddings are evaluated using pre-trained CLIP [33] f_{CLIP} . As the dependency tree examines the relations between words that could implicitly encode relational information, we use a GNN [34] to operate




Fig. 2. An overview of PARAGON. PARAGON uses soft parsing to represent language input as relations between objects. The grounding module then aligns the stated objects to the object-centric regions in the visual scene. A GNN then conducts relational reasoning between grounded objects and outputs placement. PARAGON is fully differentiable and trained end-to-end without labels for parsing and visual grounding.




Fig. 3. Dependency parsing takes language sequence as the input and outputs a tree data structure with tags. The blue blocks are dependency tags, while the red ones are part-of-speech tags. A part-of-speech tag categorizes words' correspondence with a particular part of speech, depending on the word's definition and context. Dependency tags mark two words relations in grammar, represented as Universal Dependency Relations.




Fig. 4. The pipeline of soft parsing module.

over the dependency tree from spaCy [35] for the structural features. In the dependency tree, the part-of-speech tags are node features, and the dependency tags are the edge attributes. Then, we feed the structural features into three modules to evaluate each word's weight to indicate how much this word contributes to the components in a triplet. We use a single layer MLP to compute the attention of subject words $\{a_{\text{subj},l}\}_{1 \leq l \leq L}$ and slot attention module [36] to get N features and attentions $\{a_{\text{obj},l}^n\}_{1 \leq n \leq N, 1 \leq l \leq L}$ of object phrases, where N is the total number of triplets. We feed the features of object phrases as the query, and structural features as the key and value of attention, into a multi-head attention layer [37] to get the attention for corresponding relational phrases $\{a_{\text{rel},l}^n\}_{1 \leq n \leq N, 1 \leq l \leq L}$. We compute the attention-weighted sum of word embeddings from pre-trained language models to get the embeddings of N possible relational triplets, denoted as $\Omega = \{\{\varphi_{\text{subj}}, \varphi_{\text{rel}}, \varphi_{\text{obj}}\}\}_{n=1}^N$.

Ideally, each triplet should contain information about single objects and their relations. However, we discovered that the features of multiple objects and their relations could be entangled in one triplet. As the disentangled triplets can provide more precise information and be more accurate in recovering the sentence, we add sentence recovery to guide extracting disentangled triplets. We first compute the embedding of each token using attention: $\tilde{\varphi}_l = \sum_{n=1}^N \sum_{i \in \{\text{rel}, \text{obj}\}} a_{i,l}^n \varphi_i^n + a_{\text{subj},l} \varphi_{\text{subj}}$. We then feed the embedding of tokens to an LSTM to produce the recovered word embeddings: $\tilde{l} = \{\tilde{\omega}_l\}_{l=1}^L = f_{\text{LSTM}}(\{\tilde{\varphi}_l\}_{l=1}^L)$. We minimize the L2 loss between generated sequence l and \tilde{l} as an auxiliary task to improve the soft parsing.

V. VISUAL GROUNDING

After relation extraction, we ground the mentioned objects to the visual scenes and reason about their spatial relations for placement generation. We align the visual features of objects and the embeddings of their reference expressions into the same embedding space for object grounding. We then represent relations in conditional random fields and encode them into GNNs for reasoning and placement generation.

A. Object Reference Grounding

The object reference grounding module aligns the object and subject phrase embeddings into the objects in the visual scene. Given the visual inputs, we use a pre-trained object detector (Mask-RCNN [38] or SPACE [39]) to get object bounding boxes $\{y_m\}_{m=1}^M$ and encode their visual features $\{z_m\}_{m=1}^M$ using a pre-trained CLIP [33]. We then project the object visual features \mathbf{o}_m and linguistic features in triplets φ_i , $\varphi_i \in \{\varphi_{\text{subj}}\} \cup \{\varphi_{\text{obj}}^n\}_{n=1}^N$ into the same feature space via learnable projecting matrices Φ, Ψ . We evaluate the cosine similarity $d_{\text{cos}}(\cdot, \cdot)$ with a learnable scaling factor α between the visual and text feature to get the grounding belief: $b_m^i = \text{Softmax}_m(\alpha \cdot d_{\text{cos}}(\Phi z_m^T, \Psi \varphi_i^T))$. As such, the grounding results of object i in the triplets is a set of object-centric regions with belief $\{(b_m^i, y_m)\}_{m=1}^M$.

B. Relational Reasoning for Placement Generation

1) *Spatial relations in Conditional Random Field*: The triplets essentially build up a relational graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with objects as the nodes \mathcal{V} and their spatial relations as the edges \mathcal{E} . We use the conditional random field (CRF) specified using \mathcal{G} to represent the relations between positions of the context and target objects in instruction. The variable for the vertices $\mathbf{X} = \{\mathbf{x}_i | \mathbf{x}_i \in \mathbb{R}^2\}_{i \in \mathcal{V}}$ denotes the grounded position of each context object and the placement position of the target object. We formulate the CRF as: $p(\mathbf{X} | \mathbf{o}, \Omega) \propto \prod_{(i,j) \in \mathcal{E}} \psi_{ij}(\mathbf{x}_i, \mathbf{x}_j | r_{ij}) \prod_{i \in \mathcal{V} / \{\text{subj}\}} \phi_i(\mathbf{x}_i, \mathbf{o})$. ψ_{ij} describes the spatial relations between two objects based on the spatial relation $r_{ij} \in \{\varphi_{\text{rel}}^n\}_{n=1}^N$. ϕ_i denotes the probability of the context object's position conditioned on observation, except for the target object (subj phrase). It is because the grounded position of the target object is only useful for picking rather than placing. Mean-field variational inference approximates the CRF into a mean-field $p(\mathbf{X} | \mathbf{o}, \Omega) \approx q(\mathbf{X}) = \prod_{i \in \mathcal{V}} q_i(\mathbf{x}_i)$ and optimizes the reversed KL divergence $\text{KL}(q||p)$ to converge to the multimodal distribution [40]. The equation of Mean-field

variational inference is: $\log q_i^t(\mathbf{x}_i) = c_i + \log \phi_i(\mathbf{x}_i, \mathbf{o}) + \sum_{j \in N(i)} \int_{\mathcal{X}} q_j^{t-1}(\mathbf{x}_j) \log \psi_{ij}(\mathbf{x}_i, \mathbf{x}_j | r_{ij}) \phi_j(\mathbf{x}_j, \mathbf{o}) d\mathbf{x}_j = c_i + m_{\text{obs},i}^t(\mathbf{x}_i) + \sum_{j \in N(i)} m_{ji}^t(\mathbf{x}_i)$, where $q_i^t(\mathbf{x}_i)$ receives the message $m_{ji}^t(\mathbf{x}_i)$ from its neighboring nodes $j \in N(i)$ and the message from observed node $m_{\text{obs},i}^t(\mathbf{x}_i)$, forming a message passing algorithm for $1 \leq t \leq T$. As such, the mean-field $q(\mathbf{X})$ will iteratively converge to $p(\mathbf{X} | \mathbf{o}, \Omega)$.

2) *Relational Reasoning in Mean-field GNN*: We proposed a GNN to conduct relational reasoning on the conditional random field. The GNN approximates with the mean-field message passing algorithm. To represent complex spatial distributions conditioned on high-dimensional linguistic features, we follow [17] to map the spatial variables into high-dimensional feature spaces and learn an approximate message-passing function between these variables. To handle multimodal output distributions better, we develop a method based on [22] to represent the factor of mean-field q_i as particles $\{(h_{i,k}^t, w_{i,k}^t)\}$ in message passing, where each particle encodes the position of the corresponding object.

The GNN module takes the embeddings of spatial relations $r_{ij} = \varphi_{\text{rel}}^n$ as the edge attributes to compute the message. We initialize the particles according to a normal distribution with uniform weights: $h_{i,k}^0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), w_{i,k}^0 = 1/K$. In message passing, we use a message network f_ψ to output new message by the previous inputs $h_{j,k}^{t-1}$ and edge attributes r_{ij} : $m_{ji}^t = f_\psi(h_{j,k}^{t-1}, r_{ij})$. We then compute the embedding of observation by deepset [41]: $m_{\text{obs},j}^t = f_\phi(\sum_{m=1}^M b_m^j f_{\text{pos}}(\mathbf{y}_m))$, where f_ϕ and f_{pos} are feed-forward networks. Following the form of mean-field inference, we compute the weights of the message $w_{ji,k}^t$ by the weights of particles and scores from observations computed by a weighting network g_ϕ : $w_{ji,k}^t = \eta w_{j,k}^{t-1} \log g_\phi(h_{j,k}^{t-1}, m_{\text{obs},j}^t)$.

Then, we aggregate the message from neighboring nodes and observations to compute the node features and beliefs. First, we sum up the messages and compute the new particles by an aggregate network f_u : $h_{i,k}^t = f_u(\sum_{j \in N(i)} m_{ji,k}^t + m_{\text{obs},i}^t)$. We follow the mean-field inference to evaluate new particle weights by reweighting network g_ϕ : $\log w_{i,k}^t = c + \log g_\phi(h_{i,k}^t, m_{\text{obs},i}^t) + \sum_{j \in N(i)} w_{ji,k}^t$. Next, we use resampling to avoid particle degeneracy problems. Particle degeneracy refers to the weights of all but one particle being close to zero, resulting in extremely high variance. We resample after each iteration step to build a new set of particles from the original particles: $\{(h_{i,k}^t, w_{i,k}^t)\}_{k=1}^K = \text{SoftResamp}(\{(h_{i,k}^t, w_{i,k}^t)\}_{k=1}^K)$. However, this operation is not differentiable. Soft resampling [22] makes this process differentiable via sampling from a mixture of distribution $q(j) = \alpha w_j + (1 - \alpha) \frac{1}{K}$, and the new weights are evaluated using importance sampling, resulting in new belief: $w'_k = \frac{w_j}{(\alpha w_j + (1 - \alpha) \frac{1}{K})}$. When $\alpha > 0$, soft resampling produces non-zero gradients for backpropagation. We take $\alpha = 0.5$ in training and $\alpha = 1.0$ in inference. When $t = T$, we decode the node features to get the final predictions of the target placement $\mathbf{x}_{i,k}^* = \pi_{\text{dec}}(h_{i,k}^T)$ and the weights $w_{i,k}^* = w_{i,k}^T$.

VI. END-TO-END TRAINING

We train the model end-to-end by minimizing two objectives. One is the negative log-likelihood of dataset $\bar{\mathcal{X}}$: $\mathcal{J}_l = \sum_{\bar{\mathbf{x}} \in \bar{\mathcal{X}}} -\log \sum_k w_k^* \mathcal{N}(\bar{\mathbf{x}}; \mathbf{x}_k^*, \Sigma)$, where $\bar{\mathbf{x}} \in \bar{\mathcal{X}}$ is the labeled placement and the likelihood is a mixture of Gaussian represented by each output particle of target object $\{(\mathbf{x}_k^*, w_k^*)\}_{k=1}^K$. We also minimize the auxiliary loss for improving soft parsing: $\mathcal{J}_s = \sum_{\ell \in \bar{\mathcal{X}}} \|\ell - \tilde{\ell}\|^2 = \sum_{\ell \in \bar{\mathcal{X}}} \sum_l (\omega_l - \tilde{\omega}_l)^2$. As such, the final objective is: $\mathcal{J} = \mathcal{J}_l + \lambda \mathcal{J}_s$, where λ is a hyper-parameter.

VII. EXPERIMENTS

A. Experimental Setup

1) *Tabletop Dataset*: The tabletop dataset, as shown in Fig. 5, consists of 30K visually realistic scenes generated by PyBullet [42] and NVISII [43]. The objects are sampled from 48 objects with various shapes and colors. The images contain random lighting conditions, light reflections, and partial occlusions, reflecting the challenging language grounding in the real world. We use Mask-RCNN [38] for object detection. We also select three tasks in CLIPort’s benchmark: placing objects inside a bowl or a box to test object reference grounding without positional ambiguity involved. Due to the lack of bounding box labels, we use the unsupervised object detector SPACE [39] for CLIPort’s tasks. Other tasks of CLIPort’s dataset focus on assembly or deformable object manipulation. They are not in the scope of our research.

We prepare *human-labeled instructions* to test the model with requests using human-provided natural language. However, collecting human-labeled data is expensive, while synthesized data is easily generated. Hence, we use synthesized *structured language instructions* to train the models and fine-tune the model using human-labeled data. The training dataset of structured language instructions contains 20K instructions with single spatial relations and 20K with compositions of multiple spatial relations. Our testing dataset contains 15K instructions, including instructions with unseen compositions of seen spatial relations. We also have instructions containing ambiguous reference expressions, i.e., multiple objects are semantically identical to a reference expression. We prepare 9K human-labeled instructions, where 7K instructions are for training, and the remaining 2K instructions are for testing. The human-labeled instructions are pre-collected from Mechanical Turk. We use image pairs to show the scene before and after an object is moved and let humans provide language requests for such object placement.

2) *Evaluation Metric*: We evaluate the success rate of object placement, repeated 5 times. The successful placements should satisfy all the spatial constraints given in the instructions. The placement should not be too far away from the reference objects with a threshold of 0.4 meters, which is approximated from the human-labeled dataset. The evaluated models are trained for 300K steps with a batch size of 1. We compare our performance with CLIPort [7]. Many other methods have restrictions in natural language (restricted language expressions) or visual scene (e.g., without referential ambiguity) and are not well-suited for our setting.



Fig. 5. Images in the Tabletop dataset. The dataset contains more than 40 objects shown in (a) with randomly sampled materials and colors under random lighting conditions. The scenes shown in (b) are visually realistic and reflect the challenges in the real world.

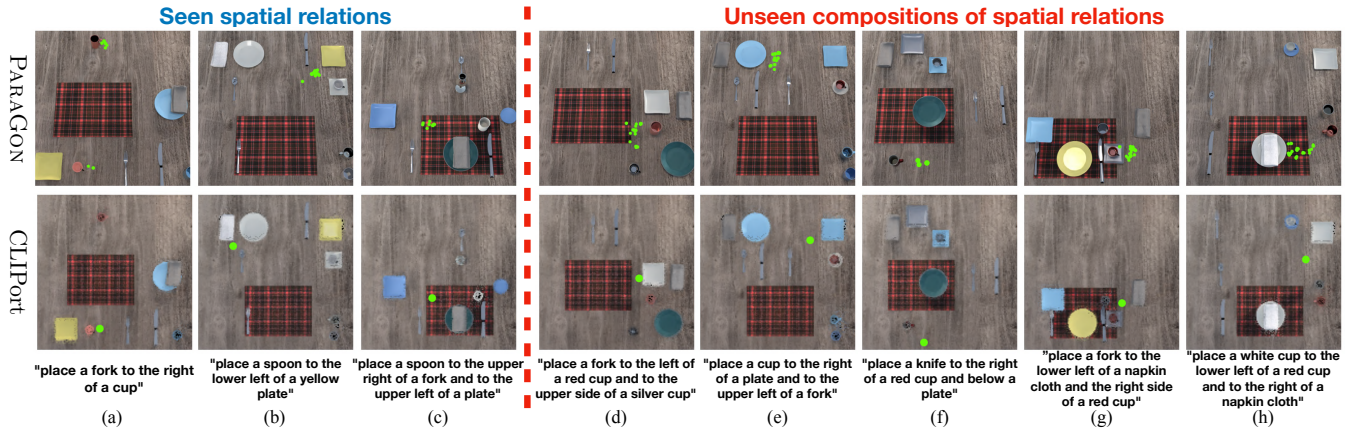


Fig. 6. Demo of placement generation by PARAGON and CLIPort in an unseen visual scene. The green dots in the first row of images are the weighted particles of placement generated by PARAGON. The green dots in the second row are argmax of placement affordance generated by CLIPort. There are demos of instructions with one spatial relation (a), seen compositions of spatial relations (b, c), unseen compositions of spatial relations (d, e, f, g, h), and ambiguous case (a). The results show that PARAGON can reason complex spatial relations for suitable placements and exhibits a sense of generalizability to unseen compositions. CLIPort hardly generates correct placement for unseen compositional instructions, indicating its poor compositional generalizability.

TABLE I
RESULTS ON CLIPORT'S DATASET: SUCCESS RATE (%)

Tasks	packing-shapes	packing-google-objects	put-blocks-in-bowls
CLIPort [7]	99.8	94.6	99.4
PARAGON	98.3	97.3	99.2

TABLE II
RESULTS ON TABLETOP DATASET: SUCCESS RATE (%)

% of Training Data	Scene Relation Type	No Ref ambiguity			With Ref Ambiguity		
		Single	Comp	Comp*	Single	Comp	Comp*
100%	CLIPort [7]	73.2	69.1	59.5	71.4	68.2	51.7
	PARAGON	93.5	92.1	90.2	93.3	91.6	89.4
10%	CLIPort	57.2	46.3	36.1	51.0	44.3	33.2
	PARAGON	89.7	88.2	88.0	89.9	87.6	87.3
2%	CLIPort	39.7	29.1	29.5	33.9	28.5	22.9
	PARAGON	86.0	69.4	64.1	85.1	67.9	65.2

B. Structured Language Instructions

We first use three tasks from CLIPort's benchmark to test the models in grounding referred objects with referential ambiguity. We then use the Tabletop dataset to test language grounding with the presence of referential ambiguity, positional ambiguity, and compositionality. We show the results on CLIPort's benchmark and Tabletop dataset in Table I and Table II, respectively.

1) *Object Reference Grounding*: In CLIPort's benchmark, spatial relations are simple, as the placements are always inside a specific object. As such, learning placement generation equals learning object grounding. In these tasks, *packing-shapes* requires object grounding without referential ambiguity,

packing-google-objects contains referential ambiguity for target objects, and *put-blocks-in-bowls* involves referential ambiguity for context objects. Table I reveals that both CLIPort and PARAGON have good performance in object grounding even with referential ambiguity.

2) *Object Placing with Ambiguity*: The Tabletop dataset involves placing the object next to a context object, causing positional ambiguity. Object grounding no longer equals object placement as the placement is not inside a specific object. Table II suggests that PARAGON has a better performance with positional ambiguity. CLIPort uses convolutional architecture, which can capture strong local correlations. This inductive bias is useful if the placement is inside an object. However, positional ambiguity is only in weak agreement with the inductive bias as pixels far away can also be closely related, compromising the performance of CLIPort. PARAGON has an encoded graphical model for the spatial relations between objects to generate placement. It learns a distribution of the usual distance for placement from data and naturally adapts to positional ambiguity.

We also test the performance of object placement with referential ambiguity (*With Ref Ambiguity* in Table II), in which the scenes contain a few semantically identical objects. Referential ambiguity makes the correct placement non-unique, resulting in a multimodal distribution of correct placement. As shown in Table II, PARAGON performs well with referential ambiguity. The core idea is to represent a distribution as a set of particles to capture multimodality and




Fig. 7. Grounding instructions with referential ambiguity results in a multimodal distribution of object placement. We use particle-based GNN to capture the multimodality to adapt to referential ambiguity.

TABLE III
RESULTS ON HUMAN INSTRUCTIONS: SUCCESS RATE (%)

Method	Fine-Tuned		Not Fine-Tuned	
	PARAGON	CLIPort	PARAGON	CLIPort
Success Rate	81.9	72.5	70.4	61.3

employ this idea in GNN for placement generation. Fig. 7 demonstrates that the GNN outputs a multimodal distribution when there is referential ambiguity.

3) *Compositionality*: Table II also reports the results on instructions with seen and unseen compositions of seen spatial relations (*Comp* and *Comp** in Table II, respectively), as well as the performance of the models trained with 2% and 10% of the training data. The composition substantially increases the complexity of language comprehension and the data required for training. Table II shows that PARAGON is data-efficient in learning structured compositional instructions and can generalize to instructions with unseen compositions. It converts a complex language into a set of simple, structured relations represented as triplets to reduce complexity. Our approach operates on the grammatical structure of natural language that is generalizable to different semantic meanings. CLIPort uses single embeddings to “memorize” seen compositions. It has poorer generalization in the presence of composition when training data is limited and does not generalize well to unseen compositions.

C. Human-labeled Instructions

We use human-labeled instructions to test natural human requests in noisy natural language. The instructions in this testing dataset may not be grammatically correct and contain unseen, noisy expressions for objects and spatial relations. The results are shown in Table III, where *Fine-Tuned* means the models are fine-tuned by the human-labeled dataset, and *Not Fine-Tuned* models are trained only on structured instructions. The results show that PARAGON performs better on human-labeled instructions. PARAGON is data-driven and optimizes all the modules to adapt to imperfect, noisy linguistic data and extract useful relational information. PARAGON extracts relations from grammatical structures of instructions, which is highly generalizable and helps tackle unseen expressions. CLIPort is also data-driven but represents instructions as single embeddings that do not generalize well to unseen, noisy language expressions.

D. Ablation Study

We conduct an ablation study to assess each module’s contribution in PARAGON. We design a variation of PARAGON without soft parsing, using single embeddings of language

TABLE IV
ABLATION STUDY: SUCCESS RATE (%)

% of Training Data	Scene Relation Type	No Ref ambiguity			With Ref Ambiguity		
		Single	Comp	Comp*	Single	Comp	Comp*
100%	PARAGON	93.5	92.1	90.2	93.3	91.6	89.4
	No Soft Parsing	73.3	69.4	33.1	80.3	71.0	36.6
	No Particle	93.1	92.3	90.5	88.3	87.4	84.1
	No Resampling	93.2	92.0	81.2	86.3	90.7	78.9
	PARAGON(ViT+Bert)	91.1	92.0	88.3	90.8	91.4	89.9
10%	PARAGON	89.7	88.2	88.0	89.9	87.6	87.3
	PARAGON(ViT+Bert)	88.4	86.0	85.2	87.7	86.5	85.4
2%	PARAGON	86.0	69.4	64.1	85.1	67.9	65.2
	PARAGON(ViT+Bert)	70.1	65.3	58.6	69.8	61.1	52.7

instructions to ground objects and represent relations. We also assess a version of PARAGON that uses a mean-field message passing neural network [17] without the use of particles, and test PARAGON without resampling. PARAGON uses CLIP, a pre-trained visual-language model, to encode text and visual inputs for object reference grounding. To test the impact of using pre-trained models that are trained separately for vision and language, we replace CLIP in PARAGON with Visual Transformer [44] and BERT [45].

We report ablation study results in Table IV. The results for *No Soft Parsing* demonstrate that soft parsing is essential in learning compositional instructions. It converts complex language sentences into simple phrases for objects and relations. The embeddings of those phrases have more straightforward semantic meanings than the entire sentence and are much easier for them to be grounded in the visual scene. As shown in the row of *No Particle*, using GNN without particles cannot capture multimodal distribution when referential ambiguity occurs and compromise the performance. The results in *No Resampling* indicate that resampling is helpful because particle degeneracy can occur without resampling, compromising the performance of capturing multimodal distributions with referential ambiguity. The last 4 rows in Table IV show that PARAGON requires fewer data to obtain good results using pre-trained CLIP than pre-trained ViT+Bert. CLIP is pre-trained by a large dataset of image-caption pairs and is better for aligning language with visual features.

VIII. CONCLUSION

PARAGON integrates parsing into a probabilistic, data-driven framework for language-conditioned object placement. It combines the strength of parsing-based and embedding-based approaches for human language grounding. PARAGON reduces the complexity of embedding-based grounding by parsing complex sentences into simple, compositional structures and learns generalizable parsing rules from data to improve robustness. Our experiments use real, noisy human instructions and photo-realistic images, reflecting the difficulties of language grounding in realistic situations to show PARAGON’s capability. For future research, grounding 3D placement such as “place a knife to lean against a bowl” is interesting. As pre-trained large language models greatly impact the field through their powerful capability of content generation, leveraging large language models to parse complex instructions also helps improve our approach.

REFERENCES

- [1] M. Shridhar and D. Hsu, "Interactive visual grounding of referring expressions for human-robot interaction," *arXiv preprint arXiv:1806.03831*, 2018.
- [2] M. Shridhar, D. Mittal, and D. Hsu, "Ingress: Interactive visual grounding of referring expressions," *The International Journal of Robotics Research*, vol. 39, no. 2-3, pp. 217–232, 2020.
- [3] O. Mees, A. Emek, J. Vertens, and W. Burgard, "Learning object placements for relational instructions by hallucinating scene representations," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 94–100.
- [4] O. Mees and W. Burgard, "Composing pick-and-place tasks by grounding language," in *International Symposium on Experimental Robotics*. Springer, 2020, pp. 491–501.
- [5] W. Liu, C. Paxton, T. Hermans, and D. Fox, "Structformer: Learning spatial structure for language-guided semantic rearrangement of novel objects," *arXiv preprint arXiv:2110.10189*, 2021.
- [6] R. Kartmann, D. Liu, and T. Asfour, "Semantic scene manipulation based on 3d spatial object relations and language instructions," in *2020 IEEE-RAS 20th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2021, pp. 306–313.
- [7] M. Shridhar, L. Manuelli, and D. Fox, "Cliport: What and where pathways for robotic manipulation," in *Proceedings of the 5th Conference on Robot Learning (CoRL)*, 2021.
- [8] S. Stepputtis, J. Campbell, M. Phielipp, S. Lee, C. Baral, and H. Ben Amor, "Language-conditioned imitation learning for robot manipulation tasks," *Advances in Neural Information Processing Systems*, vol. 33, 2020.
- [9] C. Lynch and P. Sermanet, "Language conditioned imitation learning over unstructured data," *Proceedings of Robotics: Science and Systems*. doi, vol. 10, 2021.
- [10] V. Raman, C. Lignos, C. Finucane, K. C. Lee, M. P. Marcus, and H. Kress-Gazit, "Sorry dave, i'm afraid i can't do that: Explaining unachievable robot tasks using natural language," in *Robotics: Science and Systems*. Citeseer, 2013, pp. 2–1.
- [11] A. Boteanu, T. Howard, J. Arkin, and H. Kress-Gazit, "A model for verifiable grounding and execution of complex natural language instructions," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2016, pp. 2649–2654.
- [12] T. M. Howard, S. Tellex, and N. Roy, "A natural language planner interface for mobile manipulators," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 6652–6659.
- [13] A. Boteanu, J. Arkin, S. Patki, T. Howard, and H. Kress-Gazit, "Robot-initiated specification repair through grounded language interaction," *arXiv preprint arXiv:1710.01417*, 2017.
- [14] H. Kress-Gazit, M. Lahijanian, and V. Raman, "Synthesis for robots: Guarantees and feedback for robot behavior," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 1, pp. 211–236, 2018.
- [15] S. Tellex, N. Gopalan, H. Kress-Gazit, and C. Matuszek, "Robots that use language," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 3, no. 1, 2020.
- [16] P. Karkus, X. Ma, D. Hsu, L. P. Kaelbling, W. S. Lee, and T. Lozano-Pérez, "Differentiable algorithm networks for composable robot learning," *arXiv preprint arXiv:1905.11602*, 2019.
- [17] H. Dai, B. Dai, and L. Song, "Discriminative embeddings of latent variable models for structured data," in *International conference on machine learning*. PMLR, 2016, pp. 2702–2711.
- [18] J. Dauwels, S. Korf, and H.-A. Loeliger, "Particle methods as message passing," in *2006 IEEE International Symposium on Information Theory*. IEEE, 2006, pp. 2052–2056.
- [19] F. Gustafsson, "Particle filter theory and practice with positioning applications," *IEEE Aerospace and Electronic Systems Magazine*, vol. 25, no. 7, pp. 53–82, 2010.
- [20] P. Karkus, D. Hsu, and W. S. Lee, "Particle filter networks with application to visual localization," in *Conference on robot learning*. PMLR, 2018, pp. 169–178.
- [21] M. Zhu, K. Murphy, and R. Jonschkowski, "Towards differentiable resampling," *arXiv preprint arXiv:2004.11938*, 2020.
- [22] X. Ma, P. Karkus, D. Hsu, and W. S. Lee, "Particle filter recurrent neural networks," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 04, 2020, pp. 5101–5108.
- [23] E. B. Sudderth, A. T. Ihler, W. T. Freeman, and A. S. Willsky, "Nonparametric belief propagation," in *Proceedings of the 2003 IEEE computer society conference on Computer vision and pattern recognition*, 2003, pp. 605–612.
- [24] X. Su, W. S. Lee, and Z. Zhang, "Multiplicative gaussian particle filter," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2020, pp. 56–65.
- [25] S. Schuster, R. Krishna, A. Chang, L. Fei-Fei, and C. D. Manning, "Generating semantically precise scene graphs from textual descriptions for improved image retrieval," in *Proceedings of the fourth workshop on vision and language*, 2015, pp. 70–80.
- [26] Y.-S. Wang, C. Liu, X. Zeng, and A. Yuille, "Scene graph parsing as dependency parsing," *arXiv preprint arXiv:1803.09189*, 2018.
- [27] Z. Guo, Y. Zhang, and W. Lu, "Attention guided graph convolutional networks for relation extraction," in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2019, pp. 241–251.
- [28] Q. Wang, Z. Mao, B. Wang, and L. Guo, "Knowledge graph embedding: A survey of approaches and applications," *IEEE Transactions on Knowledge and Data Engineering*, vol. 29, no. 12, pp. 2724–2743, 2017.
- [29] A. Rubinstein, "Why are certain properties of binary relations relatively more common in natural language?" *Econometrica: Journal of the Econometric Society*, pp. 343–355, 1996.
- [30] S. Kübler, R. McDonald, and J. Nivre, "Dependency parsing," *Synthesis lectures on human language technologies*, vol. 1, no. 1, pp. 1–127, 2009.
- [31] S. Petrov, D. Das, and R. McDonald, "A universal part-of-speech tagset," *arXiv preprint arXiv:1104.2086*, 2011.
- [32] R. McDonald, J. Nivre, Y. Quirnbach-Brundage, Y. Goldberg, D. Das, K. Ganchev, K. Hall, S. Petrov, H. Zhang, O. Täckström *et al.*, "Universal dependency annotation for multilingual parsing," in *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 2013, pp. 92–97.
- [33] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark *et al.*, "Learning transferable visual models from natural language supervision," in *International Conference on Machine Learning*. PMLR, 2021, pp. 8748–8763.
- [34] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl, "Neural message passing for quantum chemistry," in *International conference on machine learning*. PMLR, 2017, pp. 1263–1272.
- [35] M. Honnibal and I. Montani, "spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing," 2017, to appear.
- [36] F. Locatello, D. Weissenborn, T. Unterthiner, A. Mahendran, G. Heigold, J. Uszkoreit, A. Dosovitskiy, and T. Kipf, "Object-centric learning with slot attention," *Advances in Neural Information Processing Systems*, vol. 33, pp. 11 525–11 538, 2020.
- [37] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [38] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask r-cnn," in *2017 IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 2980–2988.
- [39] Z. Lin, Y.-F. Wu, S. V. Peri, W. Sun, G. Singh, F. Deng, J. Jiang, and S. Ahn, "Space: Unsupervised object-oriented scene representation via spatial attention and decomposition," in *International Conference on Learning Representations*, 2020. [Online]. Available: <https://openreview.net/forum?id=rkl03ySYDH>
- [40] C. M. Bishop and N. M. Nasrabadi, *Pattern recognition and machine learning*. Springer, 2006, vol. 4, no. 4.
- [41] M. Zaheer, S. Kottur, S. Ravanbakhsh, B. Póczos, R. R. Salakhutdinov, and A. J. Smola, "Deep sets," *Advances in neural information processing systems*, vol. 30, 2017.
- [42] E. Coumans and Y. Bai, "Pybullet, a python module for physics simulation for games, robotics and machine learning," <http://pybullet.org>, 2016–2021.
- [43] N. Morrical, J. Tremblay, Y. Lin, S. Tyree, S. Birchfield, V. Pascucci, and I. Wald, "Nvisii: A scriptable tool for photorealistic image generation," *arXiv preprint arXiv:2105.13962*, 2021.
- [44] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," in *International Conference on Learning Representations*, 2020.

- [45] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.